IN THE CLAIMS:

1. (Currently Amended) A method for a storage operating system implemented in a storage system to optimize the amount of readahead data retrieved for a read stream established in a data container stored in the storage system, the method comprising:

receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from the data container containing the read stream;

determining whether the storage operating system is permitted to retrieve readahead data from the data container in response to the received client read request;

if it is determined that the storage operating system is permitted to retrieve readahead data from the data container, performing the steps of:

(i) selecting an amount of readahead data to retrieve from the data container based on a plurality of factors stored within a readset data structure associated with the read stream; and

(ii) retrieving the selected amount of readahead data from the data container.

2. (Original) The method of claim 1, wherein the data container is a file, directory, vdisk or lun.

3. (Original) The method of claim 1, wherein the storage operating system is determined to be permitted to retrieve readahead data from the data container when the client-requested data extends the read stream past a predetermined next readahead value.

4. (Original) The method of claim 3, wherein the predetermined next readahead value is stored in a readset data structure associated with the read stream.

5. (Original) The method of claim 3, wherein the predetermined next readahead value is updated based on a percentage of the selected amount of readahead data.

6. (Previously Presented) The method of claim 1, wherein a read-access style associated with the data container is one of the plurality of factors used to select the amount of readahead data.

7. (Original) The method of claim 6, wherein the selected amount of readahead data equals zero if the read-access style corresponds to a random read-access style.

8. (Previously Presented) The method of claim 1, wherein a number of client read requests processed in the read stream is one of the plurality of factors used to select the amount of readahead data.

9. (Original) The method of claim 8, wherein the number of client read requests processed in the read stream is stored as a count value in a readset data structure associated with the read stream.

10. (Previously Presented) The method of claim 1, wherein the amount of client-requested data is one of the plurality of factors used to select the amount of readahead data.

11. (Original) The method of claim 10, wherein the selected amount of readahead data is set equal to a predetermined upper limit for large amounts of client-requested data.

12. (Original) The method of claim 1, wherein the selected amount of readahead data is doubled if the number of client read requests processed in the read stream is greater than a first threshold value.

13. (Original) The method of claim 1, wherein the client-requested data is identified as read-once data when either (i) the number of client read requests processed in the read stream is greater than a second threshold value or (ii) a set of metadata associated with the read stream indicates that the client-requested data is read-once data.

14. (Original) The method of claim 1, wherein the selected amount of readahead data is stored in one or more buffers enqueued on a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.

15. (Original) The method of claim 14, wherein the predetermined period of time equals two seconds.

16. (Currently Amended) An apparatus configured to implement a storage operating system that optimizes the amount of readahead data retrieved for a read stream established in a data container stored in the apparatus, the apparatus comprising:

means for receiving a client read request, the client read request indicating client-requested data for the storage operating system to retrieve from the data container containing the read stream;

means for determining whether the storage operating system is permitted to retrieve readahead data from the data container in response to the received client read request;

means for selecting an amount of readahead data to retrieve from the data container based on a plurality of factors stored within a readset data structure associated with the read stream; and

means for retrieving the selected amount of readahead data from the data container.

17. (Original) The apparatus of claim 16, wherein the data container is a file, directory, vdisk or lun.

18. (Original) The apparatus of claim 16, wherein the storage operating system is determined to be permitted to retrieve readahead data from the data container when the client-requested data extends the read stream past a predetermined next readahead value.

19. (Original) The apparatus of claim 18, further comprising means for updating the predetermined next readahead value based on a percentage of the selected amount of readahead data.

20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors used to select the amount of readahead data includes at least one of:

(i) the amount of client-requested data,

(ii) a number of client read requests processed in the read stream, and

(iii) a read-access style associated with the data container.

21. (Original) The apparatus of claim 16, wherein the selected amount of readahead data is doubled if the number of client read requests processed in the read stream is greater than a first threshold value.

22. (Currently Amended): A storage system configured to optimize the amount of readahead data retrieved for a read stream established in a data container stored in the storage system, the storage system comprising:

a network adapter for receiving a client read request, the client read request indicating client-requested data to retrieve from the data container containing the read stream; and

a memory configured to store instructions for implementing a storage operating system that performs the steps of:

determining whether the storage operating system is permitted to retrieve readahead data from the data container in response to the received client read request, and

if it is determined that the storage operating system is permitted to retrieve readahead data from the data container:

(i) selecting an amount of readahead data to retrieve from the data container based on a plurality of factors stored within a readset data structure associated with the read stream; and

(ii) retrieving the selected amount of readahead data from the data container.

23. (Original) The storage system of claim 22, wherein the data container is a file, directory, vdisk or lun.

24. (Original) The storage system of claim 22, wherein the storage operating system is determined to be permitted to retrieve readahead data from the data container when the client-requested data extends the read stream past a predetermined next readahead value.

25. (Original) The storage system of claim 24, wherein the predetermined next readahead value is updated based on a percentage of the selected amount of readahead data.

26. (Previously Presented) The storage system of claim 22, wherein the plurality of factors used to select the amount of readahead data includes at least one of:

(i) the amount of client-requested data,

(ii) a number of client read requests processed in the read stream, and

(iii) a read-access style associated with the data container.

27. (Original) The storage system of claim 22, wherein the selected amount of readahead data is doubled if the number of client read requests processed in the read stream is greater than a first threshold value.

28. (Currently Amended) A computer-readable media comprising instructions for execution in a processor for the practice of a method for a storage operating system implemented in a storage system to optimize the amount of readahead data retrieved for a read stream established in a data container stored in the storage system, the method comprising:

receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from the data container containing the read stream;

determining whether the storage operating system is permitted to retrieve readahead data from the data container in response to the received client read request;

if it is determined that the storage operating system is permitted to retrieve readahead data from the data container, performing the steps of:

(i) selecting an amount of readahead data to retrieve from the data container based on a plurality of factors stored within a readset data structure associated with the read stream; and

(ii) retrieving the selected amount of readahead data from the data container.

29. (Original) The computer-readable media of claim 28, wherein the data container is a file, directory, vdisk or lun.

30. (Previously Presented) The method of claim 1, wherein the retrieved readahead data is stored in one or more buffers, the buffers containing a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.

31. (Previously Presented) The method of claim 30, wherein the read stream corresponds to a read-once data transfer and data retrieved from the data container is stored in the flush queue.

32. (Previously Presented) The method of claim 30, wherein the retrieved readahead data is stored in the flush queue.

33. (Previously Presented) The method of claim 30, wherein one or more buffers accessed from the flush queue are re-enqueued on a normal queue.

34. (Currently Amended) A method for optimizing readahead data retrieval for a read stream established in a data container stored in a storage system, the method comprising:

receiving a client read request at the storage system, the client read request belonging to the read stream and indicating an amount of client-requested data;

selecting an amount of readahead data based on the indicated amount of client-requested data stored within a readset data structure associated with the read stream; and

retrieving the selected amount of readahead data from the data container.

35. (Previously Presented) The method of claim 34, wherein the selected amount of readahead data is set equal to a multiple of a predetermined amount, and wherein the multiple is associated with the amount of client-requested data.

36. (Previously Presented) The method of claim 34, wherein the selected amount of readahead data is set equal to a multiple of the amount of client-requested data.

37. (Previously Presented) The method of claim 36, further comprising the step of rounding the selected amount of readahead data to the size of a data block.

38. (Previously Presented) The method of claim 34, wherein the selected amount of readahead data is set equal to a predetermined upper limit.

39. (Currently Amended) A method for optimizing readahead data retrieval for a read stream established in a data container stored in a storage system, the method comprising:

receiving a client read request at the storage system, the client read request belonging to the read stream and indicating client-requested data;

selecting an amount of readahead data based on a read-access style associated with the data container, wherein the read-access style is stored within a readset data structure associated with the read stream; and

retrieving the selected amount of readahead data from the data container.

40. (Previously Presented) The method of claim 39, wherein the selected amount of readahead data equals zero if the read-access style corresponds to a random read-access style.

41. (Previously Presented) A method for optimizing readahead data retrieval for a read stream established in a data container stored in a storage system associated with a number of storage devices, the method comprising:

receiving a client read request at the storage system, the client read request belonging to the read stream and indicating client-requested data;

selecting an amount of readahead data based on the number of storage devices; and

retrieving the selected amount of readahead data from the data container.

42. (Previously Presented) The method of claim 41, wherein the step of selecting an amount of readahead data further comprises:

determining whether a flag is associated with the read stream, the flag indicating that the storage system is associated with more than a predetermined number of storage devices; and

in response to determining whether the flag is associated, selecting the amount of readahead data.

43. (Previously Presented) The method of claim 41, wherein the storage devices comprise one or more disks.

44. (Currently Amended) A method for optimizing readahead data retrieval for a read stream established in a data container stored in a storage system, the method comprising:

receiving a client read request at the storage system, the client read request belonging to the read stream and indicating client-requested data;

selecting an amount of readahead data based on a plurality of factors stored within a readset data structure associated with the read stream; and

retrieving the selected amount of readahead data from the data container.

45. (Previously Presented) The method of claim 44, wherein the retrieved readahead data is stored in one or more buffers, the buffers containing a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.

46. (Previously Presented) The method of claim 45, wherein the read stream corresponds to a read-once data transfer and data retrieved from the data container is stored in the flush queue.

47. (Previously Presented) The method of claim 45, wherein the retrieved readahead data is stored in the flush queue.

48. (Previously Presented) The method of claim 45, wherein one or more buffers accessed from the flush queue are re-enqueued on a normal queue.

49. (Currently Amended) A system for optimizing readahead data retrieval for a read stream established in a data container stored in a storage system, the system comprising:

means for receiving a client read request at the storage system, the client read request belonging to the read stream and indicating client-requested data;

means for selecting an amount of readahead data based on a plurality of factors stored within a readset data structure associated with the read stream; and

means for retrieving the selected amount of readahead data from the data container.

50. (Previously Presented) The system of claim 49, wherein the retrieved readahead data is stored in one or more buffers, the buffers containing a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.

51. (Previously Presented) The system of claim 50, wherein the read stream corresponds to a read-once data transfer and data retrieved from the data container is stored in the flush queue.

52. (Previously Presented) The system of claim 50, wherein the retrieved readahead data is stored in the flush queue.

53. (Previously Presented) The system of claim 50, wherein one or more buffers accessed from the flush queue are re-enqueued on a normal queue.

Please add new claims 54 *et al.*

54. (New) A method, comprising:

receiving a plurality of client read requests at a storage system, the client read requests indicating client-requested data sets for a storage operating system to retrieve from one or more data containers containing one or more read streams;

selecting an amount of readahead data to retrieve from the one or more data containers based on a plurality of factors stored within a readset data structure associated with each read stream;

retrieving the selected amount of readahead data from the data container;

processing one or more of the plurality of client read requests; and

adjusting, as client requests are processed, the plurality of factors stored within the readset data structure associated with each read stream to optimize amount of readahead data is cached for each read stream.

55. (New) The method of claim 54, further comprising:

determining whether the storage operating system is permitted to retrieve readahead data from the one or more data containers in response to each received client read request.

56. (New) The method of claim 54, wherein the one or more data containers are at least one of a file, a directory, a vdisk or a lun.

57. (New) The method of claim 55, wherein the storage operating system is determined to be permitted to retrieve readahead data from the one or more data containers when the client-requested data extends the read stream past a predetermined next readahead value.

58. (New) The method of claim 57, wherein the predetermined next readahead value is stored in a readset data structure associated with the read stream.

59. (New) The method of claim 57, wherein the predetermined next readahead value is updated based on a percentage of the selected amount of readahead data.

60. (New) The method of claim 54, wherein a read-access style associated with the one or more data containers is one of the plurality of factors used to select the amount of readahead data.

61. (New) The method of claim 60, wherein the selected amount of readahead data equals zero if the read-access style corresponds to a random read-access style.

62. (New) The method of claim 54, wherein a number of client read requests processed in the read stream is one of the plurality of factors used to select the amount of readahead data.

63. (New) The method of claim 62, wherein the number of client read requests processed in the read stream is stored as a count value in a readset data structure associated with the read stream.

64. (New) The method of claim 54, wherein the amount of client-requested data is one of the plurality of factors used to select the amount of readahead data.

65. (New) The method of claim 64, wherein the selected amount of readahead data is set equal to a predetermined upper limit for large amounts of client-requested data.

66. (New) The method of claim 54, wherein the selected amount of readahead data is doubled if the number of client read requests processed in the read stream is greater than a first threshold value.

67. (New) The method of claim 55, wherein the client-requested data is identified as read-once data when either (i) the number of client read requests processed in the read stream is greater than a second threshold value or (ii) a set of metadata associated with the read stream indicates that the client-requested data is read-once data.

68. (New) The method of claim 54, wherein the selected amount of readahead data is stored in one or more buffers enqueued on a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.